

## Lecture 2: Computational Security

Lecturer: Nir Bitansky

Scribes: Gal Ron, Aviv Rosenberg

## 1 Previously on Foundations of Crypto

In the previous lecture, we discussed the most basic cryptographic problem — encryption. We've defined perfect encryption and have learned that in reality it is too good to be true. We saw that if we do not want the key size  $n$  to grow with the amount of encrypted information then we have to significantly relax security:

- We have to allow a small probability that the system would break.
- We have to consider adversaries that run in bounded time.

Specifically, if the key is even slightly shorter than the message, we saw an efficient attack that simply guesses the key, and thus breaks the system with probability  $2^{-n}$ . We also saw a  $2^{n+o(n)}$ -time attack that breaks the system with very high probability. In fact, the latter attack can be sped up if we allow non-determinism.

**Claim 1.1.** *If  $P = NP$ , for any cipher for messages of length  $\ell$ , with keys of length  $n < \ell - 6$ , there exist two messages  $m_0, m_1$  and an **efficient** attacker  $A$  such that*

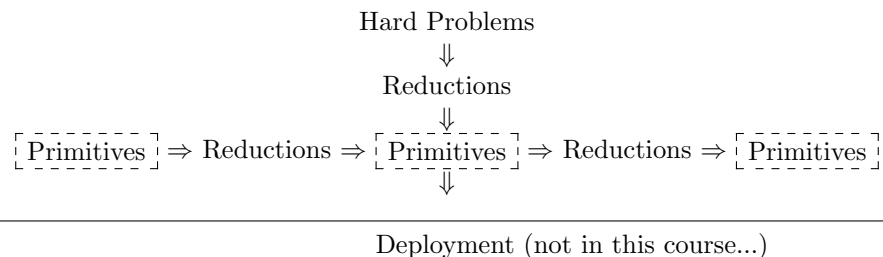
$$\Pr \left[ A(ct) = m_b \mid \begin{array}{l} sk \leftarrow \{0, 1\}^n \\ b \leftarrow \{0, 1\} \\ ct \leftarrow E_{sk}(m_b) \end{array} \right] \geq 1 - 2^{-7} > 0.99 .$$

*Proof sketch.* Recall that we proved the existence of messages  $m_0, m_1$  and a distinguisher  $A$  that simply checks if the ciphertext is in the set  $C_0 = \{E_{sk}(m_0)\}$ . This test can be easily performed in non-deterministic polynomial time (why?).  $\square$

So, computationally hard problems are necessary for cryptography. In fact, as we shall see later on, we will need much more than the assumption that  $P \neq NP$ .

## 2 The Layers of Crypto

With the above understanding, modern cryptography can be (very roughly) divided into three basic layers:



**Layer 1: Hard Problems.** These are concrete problems that we assume to be hard. Such problems may come from many different areas in different flavors:

- Number theory
- Combinatorics

- Learning
- Engineered hardness (e.g., AES)

Different assumptions may have different evidence for their hardness. Assumptions also differ in the level of confidence we have in their hardness: relations to other assumed-to-be-hard problems, etc, amount of time spent on trying to break them. Indeed, the complementing area for this layer (which we won't touch in the course but is very important) is *cryptanalysis*.

**Layer 2: Primitives.** These are cryptographic abstractions that obtain specific well-defined properties (e.g., secret-key encryption, one-way functions, fully homomorphic encryption, etc.) Such primitives are usually constructed and proven secure by a *reduction* to a hard problem. We will also try to reduce given primitives to other (hopefully, simpler) primitives.

**Layer 3: Cryptographic Systems.** Deploying cryptographic schemes in the real world to guarantee security for the user. This involves many issues and is often the bottleneck in achieving actual security:

- Verifying correctness of the implementation (naive software bugs can be disastrous).
- Integration within existing systems.
- Side channel attacks (e.g., measuring the power consumption of a device).
- Human interface pitfalls (e.g., phishing, weak passwords, etc.).

We will mostly focus on Layer 2, and sometimes touch layer 1, with two main objectives in mind:

1. Reduce cryptography to the fewest and simplest primitives possible (prove “completeness theorems”).
2. Be able to base those basic primitives on a variety of hard problems, preferably well-studied ones.

There are of course additional objectives, such as making our systems as efficient as possible, which could have a tradeoff with the strength of the computational assumptions that we make, but this generally wouldn't be our focus.

### 3 Encryption against Computationally-Bounded Adversaries

We will now define more precisely what we mean by security against computationally-bounded adversaries, staying focused on the natural problem of encryption. Throughout, our definition of the syntax and correctness of encryption schemes will remain as in the previous lecture. We focus on security.

**Definition 3.1** (Computational Security, Quantified Version). *For time bound  $t = t(n)$  and security error bound  $\varepsilon = \varepsilon(n)$ ,  $(E, D)$  is  $(t, \varepsilon)$ -secure for messages of size  $\ell = \ell(n)$  if for any two messages  $m_0, m_1 \in \{0, 1\}^\ell$ , no  $t$ -time adversary  $A$  can distinguish an encryption of one from the other with advantage greater than  $\varepsilon$ :<sup>1</sup>*

$$\Pr \left[ A(ct) = m_b \mid \begin{array}{l} sk \leftarrow \{0, 1\}^n \\ b \leftarrow \{0, 1\} \\ ct \leftarrow E_{sk}(m_b) \end{array} \right] \leq \frac{1}{2} + \varepsilon .$$

Recall that for this definition to be feasible it must be that  $t \leq 2^{n+o(n)}$  and  $\varepsilon \geq 2^{-n}$ , and ideally we'd like our encryption schemes to come as close as possible to this level of security. This is often termed *n-bit security*. Naturally, generalizing this, a system has *s-bit security*, for some  $s(n) \leq n$  if it is  $(t, \varepsilon)$  secure for  $t = \varepsilon^{-1} = 2^s$ .

<sup>1</sup>As long as we're talking about secret-key encryption, we'll stick with the convention that the secret key is simply chosen at random (see note in previous lecture).

**The Asymptotic Approach.** In practice, *every bit of security matters*. Due to efficiency considerations, the length  $n$  of keys is chosen so that  $2^{s(n)}$  is sufficiently above what is considered a feasible running time for today's strongest attackers, but not much more.

Dealing with concrete security parameters could be quite a hairy business. To understand the theoretical foundations of crypto, we don't have to go there. Instead, we will take an asymptotic approach with a quite liberal interpretation of *feasible* (or *efficient*) as *polynomial*, and infeasible as *super polynomial*. We will use the following terminology:

- We say that a function  $f(n)$  is polynomially-bounded if for some constant  $c$ , and all  $n \in \mathbb{N}$ ,  $f(n) \leq n^c$ . We sometimes denote this by  $f(n) = n^{O(1)}$ .
- We will say that an algorithm  $A$  is polynomial-time if its running time is polynomially-bounded (as a function of its input size).<sup>2</sup> (Sometimes, we will say that  $A$  is  $n^{O(1)}$ -time.) More generally, we will consider probabilistic poly-time (PPT) algorithms.
- We will say that a function  $\mu(n) \in [0, 1]$  is negligible if it decays faster than any polynomial. That is, for any constant  $c$ , there exists  $n_c$ , such that for all  $n > n_c$  it holds that  $\mu(n) \leq n^{-c}$ . (Sometimes, we will denote this by  $\mu(n) = n^{-\omega(1)}$ .)

These definitions behave quite nicely and are thus easy to work with. For instance,  $n^{O(1)} \cdot n^{O(1)} = n^{O(1)}$ ,  $(n^{O(1)})^{O(1)} = n^{O(1)}$ , and  $n^{O(1)} \cdot n^{-\omega(1)} = n^{-\omega(1)}$ .

**Pause: How to Model Polynomial-Time Computation?** One aspect that we haven't fully specified is how to exactly model adversarial algorithms. Here a natural choice is as Turing machines. Indeed, given that we only care about the running time being polynomial, all reasonable choices (e.g., RAM machines) would be equivalent. One distinction that we will make is between *uniform* algorithm and *non-uniform* ones:

- By a uniform algorithm we mean that its behavior on different input lengths is given by the same Turing machine  $A$ . In particular,  $A$  has a constant size description that doesn't scale with the input size.
- A non-uniform algorithm is specified by a sequence of algorithms  $A = \{A_n\}_{n \in \mathbb{N}}$ , one for every input length  $n$ , and the description size of  $A_n$  may (polynomially) scale with  $n$ . This for example allows modeling a realistic setting where the attacker has some auxiliary information (as part of its description) about the world (e.g., all the ciphertexts he'd seen in its lifetime). We can think w.l.o.g about any non-uniform PPT  $A$  as a family of boolean probabilistic circuits of polynomial size.<sup>3</sup> (By probabilistic, you can think that the circuits use a special gate that outputs a random bit.)

From here on in this course:

- We will consider adversarial algorithms to be non-uniform. (This may come with some loss of generality, if we also allow our reductions to be non-uniform. For now, we will not dwell on this.)
- The algorithms of any cryptographic scheme will be uniform.

We can now define an asymptotic notion of computational security. Roughly speaking, an encryption scheme is computationally secure if it is  $(t, \varepsilon)$ -secure for any polynomial  $t(n) = n^{O(1)}$  and some negligible  $\varepsilon = n^{-\omega(1)}$ . Let's be more explicit.

---

<sup>2</sup>In fact, we have already used this terminology when we required that the algorithms of which an encryption scheme consists are efficient.

<sup>3</sup>Indeed, there are several equivalent ways of modeling non-uniform algorithms. See Goldreich's book, Volume 1, Chapter 1.3.

**Definition 3.2** (Computational Security, Asymptotic Version).  $(E, D)$  is computationally secure for messages of size  $\ell(n)$  if for any non-uniform PPT adversary  $A$ , there exists a negligible function  $\mu(n)$ , such that for any  $n \in \mathbb{N}$ , and any two messages  $m_0, m_1 \in \{0, 1\}^{\ell(n)}$ :

$$\Pr \left[ A(ct) = m_b \mid \begin{array}{l} sk \leftarrow \{0, 1\}^n \\ b \leftarrow \{0, 1\} \\ ct \leftarrow E_{sk}(m_b) \end{array} \right] \leq \frac{1}{2} + \mu(n) .$$

## 4 Toward Secret Key Encryption with Short Keys

Armed with a definition, we will now start our way toward achieving it, trying to figure out what kind of hardness assumptions it requires. Where should we start? Let's start with the simplest instance of the problem that we can think of:

*Can we construct a computationally-secure encryption for messages of length  $n + 1$  with keys of length  $n$ ?*

Why start from the simplest instance? Well for once, if we cannot solve the simplest instance than we shouldn't expect to solve the general instance. Furthermore, as we shall see, looking at simple cases often helps to pinpoint the core difficulty of problems. In fact, it turns out that this is exactly the case here:

**Theorem 4.1.** *Assume there exists a computationally secure encryption scheme  $(E, D)$  for messages of size  $\ell(n) = n + 1$  (where  $n$  is the key length). Then for any polynomial  $\ell'(n)$ , there exists a computationally secure encryption scheme  $(E', D')$  for messages of size  $\ell'(n)$ .*

For now, we will give the construction, which is quite intuitive, but only a “fake analysis” to develop some intuition. Then, we will develop some concepts and tools that will allow us to complete this proof (and practically all following proofs in the course).

*(Half-Fake) Proof.* Given  $(E, D)$  we construct  $(E', D')$  as follows.

**The New Encryption**  $E'_{sk}(m)$  : given a key  $sk \in \{0, 1\}^n$  and a message  $m \in \{0, 1\}^{\ell'}$ , the algorithm works as follows:

- Let  $m = m_1 \dots m_{\ell'}$ , and let  $sk_0 = sk$ .
- For  $i = 1$  to  $\ell'$ ,
  - Sample a new secret key  $sk_i \leftarrow \{0, 1\}^n$ .
  - Sample  $ct_i = E_{sk_{i-1}}(sk_i, m_i)$ .
- Output the ciphertext  $ct = ct_1 \dots ct_{\ell'}$ .

**The New Decryption**  $D'_{sk}(ct)$  : given a key  $sk \in \{0, 1\}^n$  and a ciphertext  $ct$ , the algorithm works as follows:

- Let  $ct = ct_1 \dots ct_{\ell'}$ , and let  $sk_0 = sk$ .
- For  $i = 1$  to  $\ell'$ ,
  - Compute  $(sk_i, m_i) = D_{sk_{i-1}}(ct_i)$ .
- Output the message  $m = m_1 \dots m_{\ell'}$ .

It is not hard to see that the scheme is correct and efficient. Let's turn to security.

**The following part of the proof is fake, and is only for the sake of intuition and motivation!**

Our fake proof will rely on the intuitions that we've already developed for perfect secrecy. There, we have seen that perfect security is equivalent to saying that ciphertext distribution  $D$  for any message  $m$  is independent of  $m$ .

Let's imagine for a second that our encryption scheme  $(E, D)$  is perfect (although it's not), and show that  $(E', D')$  also is. We will "show" that for any  $m \in \{0, 1\}^{\ell'}$ :

$$E'_{sk}(m) = E_{sk_0}(sk_1, m_1)E_{sk_1}(sk_2, m_2) \dots E_{sk_{\ell'-1}}(sk_{\ell'}, m_{\ell'}) \equiv E_{sk_0}(0^{n+1})E_{sk_1}(0^{n+1}) \dots E_{sk_{\ell'-1}}(0^{n+1})$$

We will use the (fake) perfect security of  $E$  one step at a time: Without being too formal (this is a fake proof after all), at each step  $i$  we use the fact that the current secret key  $sk_i$  is independent of all other samples, and thus we can invoke the perfect security of the underlying encryption.

$$\begin{array}{ccccccc} E_{sk_0}(sk_1, m_1) & E_{sk_1}(sk_2, m_2) & \dots & E_{sk_{\ell'-1}}(sk_{\ell'}, m_{\ell'}) & \equiv \\ E_{sk_0}(0^{n+1}) & E_{sk_1}(sk_2, m_2) & \dots & E_{sk_{\ell'-1}}(sk_{\ell'}, m_{\ell'}) & \equiv \\ E_{sk_0}(0^{n+1}) & E_{sk_1}(0^{n+1}) & \dots & E_{sk_{\ell'-1}}(sk_{\ell'}, m_{\ell'}) & \equiv \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ E_{sk_0}(0^{n+1}) & E_{sk_1}(0^{n+1}) & \dots & E_{sk_{\ell'-1}}(0^{n+1}) & \equiv \end{array}$$

□

The above type of argument is generally called a *hybrid argument* — we would like to use some local guarantee on similarity of distributions to a global guarantee, by using the local guarantee many times. While the above proof is fake (in the sense that  $(E, D)$  is *not* perfectly secret), it does capture the "right" inductive intuition. We won't be able to show that each two subsequent distributions are the same, but we'll be able to show that their computationally close in some sense.

In what follows, we'll define this concept of computational closeness, which we'll heavily rely on throughout this course.

## 5 Computational Indistinguishability

Intuitively, we wish to generalize the fact that distributions are equally the same to the fact they're close in the eyes of computationally bounded adversaries. Before we go all the way and do that, it is useful to first generalize similarity to *statistical closeness*; namely, the case that distributions are not only close in the eyes of bounded observers, but also for unbounded ones.

**Definition 5.1** (Statistical Indistinguishability). *Two distributions  $X_0, X_1$  over the same support are  $\varepsilon$ -statistically-indistinguishable if for any (unbounded) distinguisher  $A$ :*

$$\Delta_A(X_0, X_1) := |\Pr[A(X_0) = 1] - \Pr[A(X_1) = 1]| \leq \varepsilon .$$

*The quantity  $\Delta(X_0, X_1) := \max_A \Delta_A(X_0, X_1)$  is the statistical distance between the distributions.*

In what sense is this the statistical distance?


**Claim 5.2.**

$$\Delta(X_0, X_1) = \frac{1}{2} \|X_0 - X_1\|_1 := \frac{1}{2} \sum_x |\Pr[X_0 = x] - \Pr[X_1 = x]| .$$

We can connect this definition to existing intuitions from previous definitions:

**Claim 5.3.** *For any  $X_0, X_1, A$ :*

$$\left| \Pr \left[ A(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow X_b \end{array} \right] - \frac{1}{2} \right| = \frac{\Delta_A(X_0, X_1)}{2} .$$



**Figure 1:** Illustration - The statistical distance between the two (continuous) distributions  $X_0, X_1$  is the surface between the graphs of the two density functions.

We can now define our notion of computational closeness, commonly called *computational indistinguishability*, which is a natural analog of the statistical indistinguishability.

**Definition 5.4** (Computational Indistinguishability, Quantified). *Two distributions  $X_0, X_1$  over the same support are  $(t, \varepsilon)$ -computationally-indistinguishable if for any  $t$ -time PPT  $A$ :*

$$\Delta_A(X_0, X_1) := |\Pr[A(X_0) = 1] - \Pr[A(X_1) = 1]| \leq \varepsilon .$$

Sticking with the asymptotic approach, we will want to consider ensembles of distributions  $X_0 = \{X_{0,n}\}_{n \in \mathbb{N}}$ ,  $X_1 = \{X_{1,n}\}_{n \in \mathbb{N}}$  that depend on some parameter  $n$  (e.g. key length).

**Definition 5.5** (Computational Indistinguishability, Asymptotic). *Two ensembles of distributions  $X_0 = \{X_{0,n}\}_{n \in \mathbb{N}}$ ,  $X_1 = \{X_{1,n}\}_{n \in \mathbb{N}}$  over the same supports are computationally indistinguishable if for any non-uniform PPT  $A = \{A_n\}_{n \in \mathbb{N}}$ , there exists a negligible  $\mu(n)$  such that for all  $n \in \mathbb{N}$ :*

$$\Delta_{A_n}(X_{0,n}, X_{1,n}) := |\Pr[A_n(X_{0,n}) = 1] - \Pr[A_n(X_{1,n}) = 1]| \leq \mu(n) .$$

We denote this by  $X_0 \approx_c X_1$ .

We can also think asymptotically about statistical distance and will say that two ensembles  $X_0, X_1$  are statistically indistinguishable if the latter definition also holds for unbounded  $A$  and denote this by  $X_0 \approx_s X_1$ .


It can be verified that this notion of computational distance has all the properties we expect it to have, such as the triangle equality, which in turn implies that indistinguishability is transitive.

**Claim 5.6.** *For any  $X_0, X_1, Y, A$ :*

$$\Delta_A(X_0, X_1) \leq \Delta_A(X_0, Y) + \Delta_A(Y, X_1) .$$

We are now ready to give a real (rather than fake) proof that our scheme  $(E', D')$  is computationally secure if the original scheme  $(E, D)$  is computationally secure. To do this, we will show that any efficient adversary  $A'$  against  $(E', D')$  can be translated into an efficient adversary  $A$  against  $(E, D)$ . This is what we call a *security reduction*; it is the bread and butter of provably-secure crypto and we'll do it all the time in this course (see figure 2).

It will be convenient to use the following indistinguishability-based definition for the computational security.



**Figure 2:** An efficient  $A'$  against  $(E', D')$  can be translated into an efficient  $A$  against  $(E, D)$ ;  $A$  translates an " $E$ -type" ciphertext into an " $E'$ -type" ciphertext.

**Definition 5.7** (Computational Security, Ind-Based Definition). *An encryption scheme  $(E, D)$  for messages of length  $\ell(n)$  is computationally secure if for any non-uniform PPT  $A = \{A_n\}_{n \in \mathbb{N}}$  there exists a negligible  $\mu(n)$ , such that for all  $n \in \mathbb{N}$ , and any two messages  $x, y \in \{0, 1\}^{\ell(n)}$ :*

$$\Delta_{A_n}(E_{sk}(x), E_{sk}(y)) \leq \mu(n) \text{ ,}$$

where  $sk \leftarrow \{0, 1\}^n$ .

It can be seen directly from Claim 5.3 that this definition is equivalent to our previous Definition 3.2 of computational secrecy.

Yet another equivalent definition:

**Definition 5.8** (Computational Security, Ensemble-Based Definition). *An encryption scheme  $(E, D)$  for messages of length  $\ell(n)$  is computationally secure if for any two sequences of messages  $x = \{x_n\}_{n \in \mathbb{N}}$  and  $y = \{y_n\}_{n \in \mathbb{N}}$  where each  $x_n$  and  $y_n$  are of length  $\ell(n)$ :*

$$\{E_{sk}(x_n) \mid sk \leftarrow \{0, 1\}^n\}_{n \in \mathbb{N}} \approx_c \{E_{sk}(y_n) \mid sk \leftarrow \{0, 1\}^n\}_{n \in \mathbb{N}} \text{ .}$$

*Making our Fake Proof Real.* Fix any non-uniform PPT  $A' = \{A'_n\}$ , and assume that it breaks  $(E', D')$ . That is, there exists a polynomial  $p$ , such that for infinitely many  $n \in \mathbb{N}$ , there exist two messages  $x_n, y_n \in \{0, 1\}^{\ell'(n)}$  such that:

$$\Delta_{A'_n}(E_{sk}(x_n), E_{sk}(y_n)) \geq 1/p(n) \text{ .}$$

From hereon, for notational simplicity, we'll omit the subscript  $n$  when it's clear from the context.

We will now define  $2(\ell' + 1)$  hybrid distributions  $H_{z,i}$  for  $z \in \{x, y\}$  and  $i \in \{0, \dots, \ell'\}$ :

$$H_{z,i} = E_{sk_0}(0^{n+1}) \dots E_{sk_{i-1}}(0^{n+1}) E_{sk_i}(sk_{i+1}, z_{i+1}) \dots E_{sk_{\ell'-1}}(sk_{\ell'}, z_{\ell'}) \text{ .}$$

Then note that  $H_{x,0} = E'_{sk}(x)$ ,  $H_{y,0} = E'_{sk}(y)$  and thus

$$\Delta_{A'}(H_{x,0}, H_{y,0}) \geq \frac{1}{p(n)} \text{ .}$$

Also note that  $H_{x,\ell'} = H_{y,\ell'} = E_{sk_0}(0^{n+1})E_{sk_1}(0^{n+1}) \dots E_{sk_{\ell'-1}}(0^{n+1})$ .

Applying the triangle inequality:

$$\sum_{z \in \{x, y\}} \sum_{i \in [\ell']} \Delta_{A'}(H_{z,i-1}, H_{z,i}) \geq \Delta_{A'}(H_{x,0}, H_{y,0}) \geq \frac{1}{p(n)} \text{ .}$$

Thus, there exists  $z \in \{x, y\}$  and  $i \in \{1, \dots, \ell'\}$  such that

$$\Delta_{A'}(H_{z,i-1}, H_{z,i}) \geq \frac{1}{2\ell' \cdot p(n)} \text{ .} \tag{1}$$

Recall that  $H_{i-1}$  and  $H_i$  only differ on the  $i$ th encryption:

- In  $H_{z,i-1}$ , we have  $E_{sk_{i-1}}(sk_i, z_i)$ ,
- In  $H_{z,i}$ , we have  $E_{sk_{i-1}}(0^{n+1})$ .

We will now construct an adversary  $A$  that breaks  $E$ . More precisely, we will show that

$$\Delta_A((E_{sk}(m, z_i), m), (E_{sk}(0^{n+1}), m)) \geq \frac{1}{2^{\ell'(n)} \cdot p(n)} ,$$

where  $sk \leftarrow \{0, 1\}^n$  is a random key and  $m \leftarrow \{0, 1\}^n$  is a random  $n$ -bit message. (Why is this enough?)

Given a ciphertext  $ct$ , and message  $m$ ,  $A$  will sample himself all the ciphertexts, except for the  $i$ th one, and plant  $ct$  as the  $i$ th cipher:

$$E_{sk_0}(0^{n+1}) \quad \dots \quad E_{sk_{i-2}}(0^{n+1}) \quad ct \quad E_{sk_i}(sk_{i+1}, z_{i+1}) \quad \dots \quad E_{sk_{\ell'-1}}(sk_{\ell'}, z_{\ell'}) .$$

It will then run  $A'$  on this sample and output the same bit. The reason that  $A$  can sample  $ct_1, \dots, ct_{i-1}, ct_i, \dots, ct_{\ell'}$  is that in both hybrids they are completely independent of  $sk_{i-1}$  and have the same distribution.

It is left to observe that:

$$\Delta_A((E_{sk}(m, z_i), m), (E_{sk}(0^{n+1}), m)) = \Delta_{A'}(H_{z,i-1}, H_{z,i}) \geq \frac{1}{2^{\ell'(n)} \cdot p(n)} .$$

Indeed, when  $ct \leftarrow E_{sk}(m, z_i)$  the emulated  $A'$  sees exactly the distribution  $H_{z,i-1}$  and when  $ct \leftarrow E_{sk}(0^{n+1})$  it sees exactly  $H_{z,i}$ .  $\square$

**What Happened to Our Previous Intuition?** Our fake proof gave us good intuition as to how to prove security — we will show that each two adjacent hybrids are close up to a negligible amount  $\mu$  and then deduce that the ends are close up to  $\mu(n) \cdot 2^{\ell'(n)}$ , which is still negligible since  $\ell'(n)$  is polynomial. We ended up giving a perhaps less intuitive proof that shows that if the end distributions are far, then there two adjacent ones that are far. Then, we deduced that there are two messages whose encryption in our underlying scheme are also far.

Why couldn't we just follow the first plan? In a nutshell, this has to do with asymptotics. While it mostly makes things simpler, it may something make our formal proofs not as natural as we expect. The problem is that unlike the perfect or even statistical case where  $\mu$ -closeness means  $\mu(n)$ -closeness *for all* attackers *for all*  $n$ . In the computational case, there is no meaning to  $\mu$ -closeness for a specific  $n$ . Indeed, there is no notion of *the best polynomial-time attacker for this  $n$* . In contrast, the notion of fairness  $\varepsilon$  for any specific attacker  $A$  and  $n$  is well defined and easy to work with.

When thinking about things informally, we should try to put asymptotics aside. It will often help to pretend that our notions are perfect/statistical, which may obey more natural intuitions.

**Note on Using Non-Uniformity in the Reduction.** In the above reduction, we've implicitly relied on non-uniformity. That is, even if  $A'$  was uniform, the  $A$  that we constructed was non-uniform — where does  $A$  get the value  $i$  for which  $A'$  has an advantage? non-uniformity! This use of non-uniformity is not really inherent. You can show for instance that if  $A$  chooses  $i$  at random, the reduction would still work. Alternatively, you can show that  $A$  can efficiently (and uniformly) find  $i$ , by sampling, and estimating the advantage for each  $i$ .